

A Theory of Indicative Bidding[†]

By DANIEL QUINT AND KENNETH HENDRICKS*

When selling a business by auction, sellers typically use indicative bids—nonbinding preliminary bids—to select a small number of bidders to conduct due diligence and submit binding offers. We show that if entry into the auction is costly, indicative bids can be informative: symmetric equilibrium exists in weakly increasing strategies, with bidders “pooling” over a finite number of bids. The equilibrium helps the seller select high value bidders with higher likelihood, although the highest value bidders are not always selected. When the number of potential bidders is large, revenue and total surplus are both higher than when entry is unrestricted. (JEL D44, D83)

Auctions with a fixed number of bidders have been the subject of a large and distinguished theoretical and empirical literature. However, in many applications, potential bidders have to acquire additional information at substantial cost before bidding. For example, in timber auctions, bidders perform “cruises” to obtain estimates of the volume and species composition of wood; in oil and gas lease sales, bidders invest in seismic surveys to learn more about the likelihood of finding hydrocarbons; and in takeover auctions, buyers conduct extensive due diligence to determine the value of the target. These costs are analogous to entry costs, in the sense that bidders are unwilling to bid without acquiring this information. The decision to incur these costs means that the number of actual bidders and the distribution of their values are endogenous. This fact has important implications for the design of optimal auctions. Since sellers bear some or all of the participation costs indirectly through lower bidder participation and bids, they have an incentive to restrict entry and select only those bidders most likely to have the highest willingness to pay. It also has important implications for implementation, which requires knowing the distribution of bidder values: when entry is costly, empirical researchers have to estimate this distribution taking into account the selection process for bidders.

This paper studies the effectiveness of using indicative bids to select bidders for an auction. This selection mechanism is commonly used in utility privatization, divestiture sales, and institutional real estate (see Kagel et al. 2008). It is used extensively in takeover auctions, which rank second only to treasury auctions in the

*Quint: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (email: dquint@ssc.wisc.edu); Hendricks: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, and NBER (email: hendrick@ssc.wisc.edu). The authors thank Serafin Grundl for excellent research assistantship, and Eric Budish, Peter Cramton, Navin Kartik, Paulo Somaini, Sheridan Titman, and attendees at numerous seminars and conferences for helpful comments. An earlier version of this paper was titled “Indicative Bidding in Auctions with Costly Entry.”

[†]Go to <https://doi.org/10.1257/mic.20160290> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

total value of assets sold each year. The value of mergers and acquisitions of US companies over the past twenty years has ranged from \$400 billion to \$1.5 trillion dollars per year;¹ about half of these deals involve an auction, and nearly all of these employ some form of indicative bidding.² Therefore, understanding how indicative bids work in takeover auctions is important in its own right. It is also important for obtaining estimates of the distribution of bidder values in these auctions and conducting counterfactual analyses of alternative selling mechanisms.

Boone and Mulherin (2009) describes how sellers control entry into takeover auctions. Sellers contact many possible bidders; those who are interested are required to sign confidentiality and standstill agreements. These agreements commit bidders not to make public their interest or bids, nor to make unsolicited bids; in exchange, they are given access to nonpublic information on the target. (Thus, in the usual auction terminology, those buyers who signed such agreements make up the set of potential bidders.) Bidders then submit preliminary indications of interest, which include an estimate or range of estimates for the price they expect to be willing to pay for the target. The bidders who report the highest willingness to pay are invited to conduct extensive due diligence and submit formal, binding bids.³ Due diligence includes access to the data room where legal and accounting teams can inspect and verify the target's contracts and financials. Most of the buyers' costs occur at this stage of the process, and the costs can run into the millions of dollars. The indicative bids themselves are costless, in that they are never paid by the buyers; and they are nonbinding because they do not restrict in any way the real offers that a buyer may subsequently make in the auction. Despite this absence of commitment, the use of indicative bids in takeover auctions suggests that sellers find them informative.

We use a two-stage model to study this practice. We assume buyers observe noisy real-valued signals about their private values, and can learn their values perfectly by incurring an entry cost. The initial signals are independently distributed, but values can be correlated conditional on the signals, allowing for the presence of common value components. Our model nests all of the models that have been considered in both the theoretical and empirical literatures on auctions with costly entry. Prior to entry, buyers are asked to simultaneously submit indicative bids or opt out. The seller commits to selecting the n buyers (typically two or three in practice) who send the highest indicative bids, with ties broken randomly. If fewer than n buyers submit indicative bids, then all of them enter the auction; if all buyers opt out, then no sale occurs. In the second stage, the selected bidders incur their entry costs, learn their values, and submit binding bids in a second-price auction. Ye (2007) shows that a

¹ 2015 Mergerstat Review (excerpt), from Business Valuation Resources, LLC, downloaded Sept. 1, 2015 from <http://www.bvresources/pdfs/2015MergerstateReviewExcerpt.pdf>.

² Using Securities and Exchange Commission (SEC) filings, Boone and Mulherin (2007) and Gentry and Stroup (2017) study over a thousand takeovers of public companies announced between 1989 and 2009. They find that in roughly half the deals, the company was offered for sale to multiple competing buyers, and that in most of these deals indicative bids were used to determine which buyers could conduct due diligence and submit binding bids.

³ Boone and Mulherin (2007) reports that in their sample of 202 auctions, the average number of buyers contacted is 21, the average number that sign agreements is 7, and the average number that submit binding bids is 1.57. Gentry and Stroup (2017) reports similar numbers for their sample. Unfortunately, neither reports the number of buyers submitting indicative bids.

fully separating equilibrium fails to exist in this model. We are interested in characterizing the symmetric equilibria that do exist.

Our central result is that indicative bids yield a partial sorting of buyers based on their signals (i.e., types) when the expected rents from the private information obtained in the second stage are small relative to entry costs. This condition is satisfied when most of the learning occurs prior to entry, or when the information obtained in the second stage is highly correlated across bidders. In this case, the buyers' incentives are sufficiently aligned with those of the seller that indicative bids are informative: the seller wants to restrict the number of buyers who incur the entry cost, and buyers want to avoid being selected and paying this cost if they are unlikely to win. Low-value buyers will try to separate themselves from high value buyers by submitting lower indicative bids. We show that a symmetric equilibrium is a finite partition of the space of buyers' types. Buyers with types in the same element of the partition submit the same indicative bid, and buyers in higher elements submit higher bids. Thus, the equilibrium helps the seller select high value buyers with greater likelihood. We prove existence of a symmetric equilibrium (uniqueness can be shown in some special cases), and explore some comparative statics.

How well does the indicative bidding mechanism perform? A natural benchmark is an auction in which entry is unrestricted. Buyers decide on the basis of their private information whether or not to enter the auction, pay the entry cost, update their values, and submit binding bids. Our main theoretical result is that indicative bidding yields greater revenue and greater total surplus than the unrestricted auction when the number of potential buyers is large.⁴ Under an additional assumption, ex ante bidder surplus is higher with indicative bidding as well. Through numerical examples, we find that these results tend to hold even when the number of buyers is small. Thus, when entry is costly, the introduction of indicative bids does not involve the standard trade-off in optimal auctions between revenues and efficiency.

Indicative bidding does better than unrestricted entry but, as Lu and Ye (2017) points out, it is hardly optimal within the class of two-stage mechanisms. The main alternative proposed in the literature is an entry rights auction (Fullerton and McAfee 1999, Ye 2007). In characterizing the optimal two-stage mechanism, Lu and Ye (2017) shows that it can be implemented under certain conditions using an all-pay auction for entry rights followed by a second-price auction with handicaps. However, in our view, an entry rights auction may not always be feasible, particularly in the case of takeover auctions, because it requires buyers to commit to paying substantial sums *before* they conduct due diligence. This would undermine the purpose of due diligence and put management at risk of shareholder lawsuits if the asset turned out to be worth less than they anticipated. Even worse, since revenue is based in part on the expected, rather than actual, valuation of the asset, there is a risk of adverse selection among sellers: a rush of entry by sellers with worthless assets could crowd out sellers with legitimate ones and increase the need for due diligence. If mechanisms requiring payments before bidders perform due diligence are ruled

⁴We compare auctions without reserve prices, but in the unrestricted case, the optimal reserve price goes to zero as the number of buyers gets large; so when this number is large enough, indicative bidding outperforms the optimal standard auction.

out, then it is not clear that one can do much better than indicative bidding followed by an auction.⁵

This paper contributes to the large literature on auctions with costly entry. The focus of the theoretical literature is on characterizing rational entry and bidding decisions for a given entry process and information environment, and on designing mechanisms that are more efficient or generate more revenue. One branch studies environments in which buyers have no private information prior to entry and pay a cost to learn their value. In this setting, the seller does not have to worry about selecting buyers, which is the central issue of our paper. Levin and Smith (1994) characterizes the mixed strategy entry equilibrium in private value English auctions, and shows that the failure of bidders to coordinate entry leads to outcomes that could be improved by capping entry at a fixed number. Crémer et al. (2009) characterizes the optimal mechanism, and shows that the seller can use entry fees and subsidies to extract all buyers' surplus. In the absence of such payments, Bulow and Klemperer (2009) demonstrates that auctions with unrestricted entry are less efficient than sequential mechanisms but typically generate more revenue.⁶ A second branch studies environments in which buyers know their value, or have a signal of their value, before paying an entry cost. The entry cost can be bidding costs as in Samuelson (1985) or additional information acquisition costs as in Ye (2007). This setting leads to selective entry, since bidders with high values are more likely to enter than bidders with low values. Samuelson (1985) shows that the entry equilibrium has a threshold property and that restricting the number of bidders could increase revenues. Ye (2007) shows that an entry rights auction can be used to induce efficient entry. Lu and Ye (2017) characterizes the optimal two-stage mechanism. Our contribution to this literature is to model the use of "cheap talk" to resolve the coordination problem faced by buyers.

Recently, an empirical literature on auctions with costly entry has developed. The focus of this literature is on identifying and estimating the joint distribution of bidder signals (or entry costs) and values in order to evaluate different entry/auction formats. Athey, Coey, and Levin (2013) and Athey, Levin, and Seira (2011) estimate a model in which loggers and mills in timber auctions pay an entry cost to learn their values. From an econometric perspective, the key simplification is that entry is not selective: costs and values are independently distributed. Bhattacharya et al. (2014), Li and Zheng (2009, 2012), Roberts (2013), and Roberts and Sweeting (2013) extend this model to account for selective entry by estimating parametric models of bidding for highway contracts and timber. In their models, bidders receive signals affiliated with their values before paying the cost to learn their values, and the (signal, value) pairs are independently distributed across bidders. Marmer et al. (2013) provides nonparametric tests to distinguish between the two kinds of entry models; Gentry and Li (2014) provides conditions under which the joint distribution of signals and values is identified. In a recent working paper, Gentry and

⁵One could still improve on the mechanism we study by varying the number of bidders n advancing to the auction in response to the indicative bids received, or (as we show in one of our theorems) by partly subsidizing entry when the number of potential bidders is large.

⁶Other papers in this branch include McAfee and McMillan (1987), Burguet and Sákovic (1996), Menezes and Monteiro (2000), Tan (1992), Ye (2004), and Compte and Jehiel (2007).

Stroup (2017) apply the affiliated signal model to estimate the joint distribution of signals and values in takeover auctions from data on the numbers of potential and actual bidders and the deal premium. They use their estimates to examine whether takeover auctions generate higher prices than negotiations. On the normative side, our contribution to this literature is to provide an alternative entry format that should be considered, at least for auctions that meet our small rents condition (which is a testable restriction). On the positive side, our analysis informs empirical researchers studying takeover auctions as to how they can use the number (and values) of indicative bids to help identify the joint distribution of signals and values.

Our paper also contributes to the voluminous literature on “cheap talk” games. Our indicative bidding equilibria are similar to the “cheap talk” equilibria of Crawford and Sobel (1982): indicative bids are monotonic in buyers’ initial information, but only a finite number of different bids are used in equilibrium, and different types of buyers “pool” on the same bid. In their seminal paper, Crawford and Sobel show that cheap talk can improve the ex ante payoffs of both parties when a biased sender has information relevant to the receiver’s decision problem. Farrell and Gibbons (1989) and Matthews and Postlewaite (1989) similarly show that cheap talk can be informative prior to bilateral bargaining, and can therefore expand the set of equilibrium payoffs. Our contribution to this literature is to introduce a natural kind of receiver commitment into a cheap talk setting, which sharpens the predictions of the model. As Farrell and Gibbons (1989) have observed, in standard cheap talk games, the receiver cannot commit to a choice of outcome as a function of the messages. Instead, the messages derive meaning only from the receiver’s interpretation of them and the receiver must act optimally given that interpretation. In our setting, we assume the seller commits both to the rules of the auction (which is standard) *and* to how he will select entrants based on the indicative bids received. In particular, we assume the seller commits to selecting a fixed number of bidders, choosing those who send the highest indicative bids, and breaking ties randomly. This commitment to a monotone selection rule eliminates much of the multiplicity of equilibria that arises in cheap talk games. In particular, it rules out a “babbling” equilibrium, and any equilibrium where adverse off-equilibrium-path beliefs are used to deter unused messages.⁷

We should also mention two other related papers that feature entry models whose equilibria have similar structures to ours, but where voluntary delay rather than communication is used to screen and coordinate entrants. Levin and Peck (2003) does this in an oligopoly setting, where post-entry competition is symmetric and firms have private information about their entry costs; with multiple discrete opportunities to enter, a firm with intermediate-level costs will wait to see that his opponent does not enter for a certain number of periods before entering himself. McAdams (2015) considers second-price auctions with costly bidding and reserve prices, where bids can be submitted in discrete rounds and are made public after each round. In equilibrium, a bid deters any future bids, and a new set of bidder types become

⁷Navin Kartik and Joel Sobel (private communication, August 27, 2015) have similarly shown that imposing monotonicity on both the sender’s and receiver’s strategies in a standard cheap talk setting, combined with iterated weak dominance, uniquely selects the “most informative” cheap talk equilibrium.

willing to bid (if nobody has bid prior) in each round. Like our model, both of these models admit a unique symmetric equilibrium in thresholds.

The paper proceeds as follows. Section I presents the model. Section II characterizes symmetric equilibrium and proves equilibrium existence. In Section III, we use an example to illustrate the construction of the symmetric equilibrium. In Section IV, we evaluate the performance of the indicative bidding mechanism against the alternative of unrestricted entry, and establish general results for the case where N is large. Section V discusses extensions. Section VI concludes. Proofs are in the Appendix; some lengthier proofs are in a separate online Appendix.

I. Model

We begin by describing the environment and the indicative bidding mechanism we consider and the assumptions we make.

Our model is based on the “private value updating” model of Lu and Ye (2017). There are $N \geq 3$ potential bidders, indexed by i . Bidder i 's value of the asset is V_i . Initially, she does not know V_i , but observes a real-valued, private signal S_i of it, which we will refer to as her type. She learns V_i perfectly during due diligence. We assume that all buyers face the same cost of performing due diligence, denote this cost by c , and also refer to it as the entry cost.⁸

Let $V = (V_1, \dots, V_N)$ and $S = (S_1, \dots, S_N)$. (Throughout, we use capital letters to denote random variables, and the corresponding lowercase letters to indicate their realizations.) We assume the joint distribution of (V, S) is continuous and exchangeable with respect to bidder indices. We assume that $\{S_i\}$ are independent across i , and that S_i is independent of $\{V_j\}_{j \neq i}$, but we allow for the possibility that $\{V_i\}$ are correlated conditional on S , i.e., that the new information learned during due diligence is correlated across bidders. (The literature on optimal mechanism design in this setting makes the assumption that (V_i, S_i) are independent across i , to rule out full surplus extraction à la Crémer and McLean (1988). While we will maintain the assumption of independent types $\{S_i\}$, since we study a particular mechanism, we have no need to assume that $\{V_i\}$ are independent given S .⁹) We assume that S_i has finite support $[\underline{S}, \bar{S}]$, and a continuous (marginal) distribution that admits a continuous density bounded below on its support. Under these assumptions, there is no loss of generality in assuming that S_i is distributed uniformly on $[0, 1]$, which we therefore assume.¹⁰

⁸If the seller also incurs costs directly for each buyer who goes through this process, it would strengthen the seller's incentive to limit entry, but would not change bidder play within the mechanism we study.

⁹Once (S_i, V_i) are assumed to be independent across i , one can use the normalization from Eső and Szentes (2007) and, without loss of generality, express V_i as $u(S_i, T_i)$, where u is increasing in both arguments, $\{T_i\}$ are independent across i , and T_i is independent of S_i and uniformly distributed on $[0, 1]$. The interpretation is that T_i is the new information that buyer i obtains about her value during due diligence. Our model is more general than this, primarily by allowing $\{T_i\}$ to be correlated across bidders. In addition, however, when $\{V_i\}$ are not independent, the expression of V_i as $u(S_i, T_i)$, with T_i independent of S_i , is no longer a normalization but a substantive assumption, which we have no need to impose.

¹⁰That is, if F_S denotes the marginal distribution of S_i , one could equivalently assume bidder i observed $F_S(S_i)$, which contains the same information as S_i but is distributed uniformly on $[0, 1]$.

The indicative bidding mechanism we consider consists of a cheap talk stage and an auction stage. In the cheap talk stage, the seller asks potential buyers if they are interested in bidding for the asset, and if so, how much they are willing to pay. These bids are not binding, and are known as indicative bids. An indicative bid can be a precise number, although often the seller asks the bidder to report a range in which she believes her willingness to pay is likely to fall. We formalize this stage of the game by assuming that buyers simultaneously send messages to the seller. A message is denoted by m , and the set of messages available to each bidder is the set of integers $\{0, 1, \dots, \bar{M}\}$, where “0” is “opt out,” or decline to participate, and \bar{M} is allowed to be either finite or infinite. Note that \bar{M} is a parameter of choice for the seller: for example, the seller could simply ask each buyer to opt in or opt out, in which case $\bar{M} = 1$. The substantive restriction here is that the set of opt-in messages is bounded below (i.e., there is a lowest “opt-in” message), fully ordered, and countable. Given this restriction, there is no loss of generality in assuming that the message space consists of nonnegative integers.¹¹

The auction stage consists of an English auction. The seller selects bidders for the auction based on the messages received in the cheap talk stage. We assume the seller commits to a maximal number of bidders n and, if more than n bidders opt in, commits to selecting the bidders who sent the highest messages, breaking ties randomly. If all bidders opt out, then the game ends with no sale. If only one bidder opts in, then that bidder gets the asset at the reserve price. (For expositional and notational ease, we focus on the case where the reserve price is 0 and a lone entrant therefore acquires the asset for free; a positive reserve is easily incorporated into the model and does not change our results.) In addition, the seller commits not to make public the messages that the bidders send to him, nor his response to those messages. Therefore, each bidder knows only whether or not she has advanced to the auction, and does not gain any additional information about which (if any) of the other bidders have also advanced. This nondisclosure commitment is important because if a bidder were to learn that she is facing maximal competition in the auction, it would sometimes be in her interest to drop out of the bidding to avoid incurring the entry cost.¹²

In an English auction, a bidder who advances has a (weakly) dominant strategy to bid her valuation, regardless of the message she sent in the first stage. Therefore, in what follows, we assume that conditional on advancing, bidder i bids V_i in the auction. Let $S_{-i,n}$ and $V_{-i,n}$ denote, respectively, the vectors of signals and values of bidder i 's $n - 1$ opponents in the auction. Then, conditional on being selected and performing due diligence, bidder i 's ex post payoff before costs is V_i if she advances alone, and $\max\{0, V_i - \max\{V_{-i,n}\}\}$ if she advances against $n - 1$ opponents. Given S_i , her interim expected payoff before costs from advancing against $n - 1$ opponents with types $S_{-i,n}$ is therefore

$$u_n(S_i, S_{-i,n}) \equiv E_{(V_1, \dots, V_n | S_1, \dots, S_n)}(\max\{0, V_i - \max\{V_{-i,n}\}\})$$

¹¹ As we discuss in the Appendix, symmetric equilibria fail to exist when the set of allowed messages is continuous, or more accurately, when for any two permitted messages, there is always another message between them.

¹² More generally, nondisclosure strengthens a seller's bargaining position when only one bidder is seriously interested in buying the asset. Subramanian (2010) offers a couple of funny stories on how sellers keep bidders in the dark as to the number of competitors they face.

and her expected payoff is

$$u_1(S_i) \equiv E_{V_i|S_i}(V_i)$$

if she advances alone.

We assume that expected auction payoffs depend on initial types in the expected way (or that types are ordered in the natural direction).

ASSUMPTION 1:

- (a) $u_1(S_i)$ is continuous and strictly increasing in S_i .
- (b) For each $n > 1$, $u_n(S_i, S_{-i,n})$ is continuous in its n arguments; weakly increasing in S_i ; weakly decreasing in $S_{-i,n}$; and strictly increasing in S_i if $S_i \geq \max\{S_{-i,n}\}$.

If $\{S_i, V_i\}$ are independent across i , then a sufficient condition for the monotonicity conditions to hold is for (S_i, V_i) to be affiliated. (The continuity conditions follow from continuity of the joint distribution of (S, V) .) Finally, we assume that the entry cost c is neither “too large” nor “too small.”

ASSUMPTION 2:

- (a) $c < u_1(1)$.
- (b) For every $s_i \in [0, 1]$, $u_2(s_i, s_i) < c$.

Assumption 2(a) is simply the requirement that the game is nontrivial: that the entry cost is not so large as to completely preclude entry. Assumption 2(b) is more substantive: it requires the entry cost to be high enough that a bidder only wants to enter the auction against an opponent if her type (i.e., initial signal) is greater than her opponent’s type. It is essentially a restriction on the information rents bidders can earn from due diligence, and we therefore refer to it as the “small rents” assumption.

While Assumption 2(b) is most transparently satisfied when little is learned during due diligence, it does not require this; what it requires is that the information different bidders learn about their valuations during due diligence is highly correlated. For example, if $V_i = S_i + T_i$, with $\{T_i\}$ independent of $\{S_i\}$ representing a new signal each bidder observes during due diligence, then Assumption 2(b) would be satisfied if $\{T_i\}$ were sufficiently highly correlated across i , since bidders would compete away any rents associated with learning the realization of T_i . In the extreme case where $\{T_i\}$ are perfectly correlated, we could think of them as a common component of value, with bidders already knowing the “idiosyncratic” part of their valuation for the company (synergies across activities, for example). This may be a reasonable approximation for corporate acquisitions, since due diligence is mostly about checking for hidden liabilities and other “skeletons in the closet,” which might

impact all buyers in a similar way. (Ultimately, whether the information rents earned during due diligence are large or small relative to participation costs is an empirical question, and we plan to address it in future empirical work.)

Combined with our other assumptions, Assumption 2(b) implies that there is some $\varepsilon > 0$ such that $u_2(s_i, s_j) \geq c$ requires $s_i \geq s_j + \varepsilon$. This gives some intuition for Ye's (2007) result that an indicative bidding game cannot have a fully separating equilibrium: in any symmetric equilibrium, if her opponents were bidding "truthfully," a bidder would have an incentive to misrepresent her type downward by at least ε , to avoid being selected in some scenarios where her expected payoff in the auction would be less than the participation cost. Thus, this ε is in a sense analogous to the bias between sender and receiver preferences in Crawford and Sobel (1982).¹³ However, the fact that buyers do not always gain from participating in the auction is what ensures that buyer and seller incentives are sufficiently aligned for cheap talk to be informative. The seller wants to restrict the number of buyers (since he bears some of the participation cost indirectly through its effect on participation), without excluding the strongest buyers; while the buyers want to avoid being selected and paying the entry cost when they are unlikely to win. Thus, low types will try to separate themselves from high types by sending lower messages; and sellers will happily exclude them in favor of bidders with higher types. (In contrast, if too much idiosyncratic learning occurred during due diligence, then buyers would always want to enter the auction regardless of their types, and cheap talk would unravel. We discuss relaxing Assumption 2(b) in a later section.)

The indicative bidding game can be thought of as a cheap talk game with commitment. The messages of the bidders (senders) influence which action the seller (receiver) takes but, given that action, they do not affect the payoffs of the players. In the standard cheap talk game, the space of messages is unrestricted, and the receiver chooses an action that is his best response to the messages sent. By contrast, in the indicative bidding game (as is standard in auctions), the seller commits to the mechanism. In particular, he commits to a rule that selects bidders based on the messages they send, and ignores messages outside the set of allowed ones.¹⁴

II. Equilibrium Characterization and Existence

Next, we define strategies, payoffs, and equilibrium for this model, establish some properties of any symmetric equilibrium, and establish that a symmetric equilibrium always exists.

¹³We thank Navin Kartik for this observation. Note that the "bias" in our model goes in the opposite direction as in Crawford and Sobel: in the classic sender-receiver game, the sender wants to misrepresent his type upward, while in our setting, a buyer would like to misrepresent her type downward.

¹⁴Ex interim, once the messages are received, it would typically be in the seller's interest to allow more than n bidders to advance. However, on the equilibrium path, subject to the constraint of not advancing more than n bidders, the seller is selecting the bidders he would want to if he was not committed to a selection rule. Thus, as we discuss in Section V, the equilibrium we find in the next section would still be an equilibrium if the seller had complete freedom to choose whichever bidders he wanted, rather than being committed to choosing those who submitted the highest indicative bids; but like in standard cheap talk games, we would also have less-informative equilibria in this case.

Given our model, the particular environment facing a seller is characterized by the number of bidders N , the joint distribution of signals and valuations (S, V) , and the entry cost c . The particular indicative bidding mechanism faced by bidders is characterized by the maximal number of bidders who will advance to the auction n and the number of opt-in messages allowed \overline{M} . A pure strategy for bidder i consists of a message function

$$\tau_i : [0, 1] \rightarrow \{0, 1, \dots, \overline{M}\}$$

that maps the set of types to the set of messages; a mixed strategy is a mapping from $[0, 1]$ to probability distributions over $\{0, 1, \dots, \overline{M}\}$. The support of a strategy is defined to be the set of messages played with positive probability by a positive measure of types.

Our objective is to characterize symmetric equilibria. Consequently, we need only define the expected payoff to bidder i when her $N - 1$ opponents all play a common strategy. Let $v_\tau(m; s)$ denote the expected payoff to a bidder with type s if she sends message m and her opponents are all playing the strategy τ . (Since the game is symmetric, we do not index $v_\tau(\cdot; \cdot)$ by the identity of the bidder i .) A pure-strategy symmetric Bayesian-Nash equilibrium (BNE) is then a strategy τ such that $v_\tau(\tau(s); s) \geq v_\tau(m'; s)$ for all $m' \in \{0, 1, \dots, \overline{M}\}$ and $s \in [0, 1]$.¹⁵

Under the assumptions stated above, we can establish two key properties which must hold in any symmetric equilibrium, and which allow us to characterize what any symmetric equilibrium must therefore “look like.”

LEMMA 1: *Given an environment and an indicative bidding mechanism, if τ is a symmetric equilibrium, then (i) τ is weakly increasing, and (ii) τ has support $\{0, 1, \dots, M\}$ for some finite M .*

Thus, any symmetric equilibrium (should one exist) has the same structure as the cheap talk equilibria of Crawford and Sobel (1982): the type space $[0, 1]$ is partitioned into a finite number of subintervals $[0, \alpha_0]$, $[\alpha_0, \alpha_1]$, \dots , $[\alpha_{M-1}, 1]$; bidders with types in the interior of each subinterval pool on the same message; and bidders with types at the boundary α_m between two subintervals send either message m or $m + 1$ or mix between the two.¹⁶

Lemma 1 says that even when \overline{M} is infinite, only finitely many messages are used in equilibrium. Since we assume the seller is committed to advancing the bidders who sent the highest messages, it may seem counterintuitive that even a bidder with the highest type is not willing to separate herself by sending a higher, out-of-equilibrium message and advancing for certain. However, were she to make such a deviation, the increase in her probability of advancing would arise solely from breaking ties against opponents who are sending the highest equilibrium

¹⁵ A mixed-strategy symmetric BNE is a strategy τ , such that $v_\tau(m; s) \geq v_\tau(m'; s)$ for all $m' \in \{0, 1, \dots, \overline{M}\}$, all m on which $\tau(s)$ places positive probability, and all $s \in [0, 1]$.

¹⁶ In the knife-edge case of indifference, a bidder with type $S_i = 1$ might send any message in the set $\{M, M + 1, \dots, \overline{M}\}$, or any mixture among them.

message M , and therefore have types within the highest subinterval $[\alpha_{M-1}, 1]$. As more messages get used, this interval gets sufficiently small that under Assumption 2(b), even the highest type's payoff against a randomly selected opponent drawn from this interval would be negative. When this is the case, she will not want to separate herself by sending a higher message.

Lemma 1 gives only necessary, not sufficient, conditions for a strategy τ to be a symmetric equilibrium. It turns out that in addition to these conditions, two more turn out to be sufficient.

LEMMA 2: *Let τ be a weakly increasing strategy with support $\{0, 1, \dots, M\}$ for some $M < \infty$. For $m = 0, 1, \dots, M - 1$, define α_m as the supremum of the set of bidder types who send message m with positive probability. Then τ is a symmetric equilibrium if and only if:*

- (i) $v_\tau(m; \alpha_m) = v_\tau(m + 1; \alpha_m)$ for $m = 0, 1, \dots, M - 1$, and
- (ii) either $M = \bar{M}$ (the support of τ includes all available messages) or $v_\tau(M; 1) \geq v_\tau(M + 1; 1)$, and $\tau(1)$ puts probability 1 on message M unless $v_\tau(M; 1) = v_\tau(M + 1; 1)$.

Necessity of both of these conditions is straightforward: the first follows from continuity of each payoff function $v_\tau(m; \cdot)$, and the second is required for a bidder with type $S_i = 1$ to be playing a best-response. The significant part of Lemma 2 is that these conditions are sufficient: that if a partitioned strategy is found satisfying indifference of the threshold types, and unused messages (if they exist) are not a profitable deviation for a bidder with the highest type, then the strategy in question is a symmetric equilibrium. This allows us to prove existence constructively, via an algorithm that finds exactly such a strategy, leading to the following result.

THEOREM 1: *Fix an environment and an indicative bidding mechanism. A symmetric equilibrium exists. Further, given the environment and n , there is a number M^* , with $1 \leq M^* < \infty$, such that ...*

- if $\bar{M} \leq M^*$, a symmetric equilibrium exists in which all available messages are used with positive probability;
- if $\bar{M} > M^*$, a symmetric equilibrium exists in which only the messages $\{0, 1, \dots, M^*\}$ are used.

Thus, a symmetric equilibrium can always be found in which exactly $\min\{\bar{M}, M^*\}$ opt-in messages are used.

In two special cases, we can show that this is essentially the *unique* symmetric equilibrium. Maintaining the normalization that $S_i \sim U[0, 1]$, if either ...

- (i) $V_i = u(S_i)$, with u increasing and weakly convex, or
- (ii) $n = 2$ and $V_i = \beta S_i + T_i$, with $\{T_i\}$ independent of $\{S_i\}$,

then the equilibrium found in Theorem 1 is (up to the strategies of the indifferent types $\{\alpha_m\}$) the *only* symmetric equilibrium.¹⁷ In the more general case, we are not certain whether uniqueness holds, and if not, how much multiplicity to expect. However, in a sense, this can be thought of as an empirical problem. The constructive proof of Theorem 1 also offers a simple and computationally feasible way, given a particular environment and mechanism, to exhaustively search for *all* symmetric equilibria.

III. Illustration of Equilibrium Construction

A. Preliminaries

In this section, we use an example to illustrate the construction of a symmetric equilibrium. For the example, we let $V_i = S_i$, and therefore assume that $\{V_i\}$ are independently and identically distributed uniformly on $[0, 1]$ and known perfectly prior to due diligence.¹⁸ In this case, the payoffs to a bidder from advancing to the auction stage depends only on the number and distribution of opponents who send the *highest* message, as opponents who send lower messages have lower values and are certain to bid less and lose. This property helps to simplify the calculation of bidders' payoffs.

Suppose bidder i with type s_i is told that she has been selected and that her k highest opponents have types drawn (uniformly) from the interval $[a, b]$. For a given realization s^* of the highest opponent's type, bidder i 's payoff is $\max\{0, s_i - s^*\} - c$. The CDF of this highest opponent type is $\Pr(s^* < s) = ((s - a)/(b - a))^k$, and its density function is therefore $k(s - a)^{k-1}/(b - a)^k$. Thus, conditional on her own type s_i , advancing to the auction, and facing a set of opponents of whom k have types in $[a, b]$ and the remainder have types below a , we can write bidder i 's expected payoff as

$$V(s_i, k, [a, b]) = \int_a^b \max\{0, s_i - s\} \frac{k}{b - a} \left(\frac{s - a}{b - a}\right)^{k-1} ds - c.$$

Note that this payoff does not depend on the message bidder i herself sent, nor on the messages of opponents who did not advance.

As noted earlier, for any strategy τ satisfying the conditions of Lemma 1, we can associate with the strategy a series of thresholds $\alpha_0 < \alpha_1 < \dots < \alpha_{M-1}$ separating the subintervals of types sending each message. Since $\alpha = (\alpha_0, \dots, \alpha_{M-1})$ fully determines the strategy of all but a measure 0 of bidder types, a bidder's expected payoff depends on her opponents' strategy τ only through the thresholds α associated with it; hereafter, we will write expected payoffs as $v_\alpha(\cdot; \cdot)$ rather

¹⁷In these two cases, the additional properties (discussed in the online Appendix) that allow us to establish uniqueness of the symmetric equilibrium also lead to proofs of certain comparative statics. For example, in both cases, M^* is decreasing in c ; in the latter case, M^* is also decreasing in N , and (if $\{T_i\}$ are independent across bidders) weakly increases when a mean-preserving spread is applied to the distribution of T_i .

¹⁸This is also equivalent to the case where $V_i = S_i + T_i$, where $\{T_i\}$ are learned during due diligence but are perfectly correlated across bidders, so that if two or more bidders enter, the rents from the realization of $\{T_i\}$ are fully competed away.

than $v_\tau(\cdot; \cdot)$ to emphasize this. For any opt-in message, the *unconditional* expected payoff to bidder i if her type is s_i , she sends message m , and her opponents are all playing a strategy τ described by a series of thresholds α is given by

$$\begin{aligned}
 (1) \quad v_\alpha(m; s_i) = & \alpha_0^{N-1} (s_i - c) \\
 & + \sum_{m'=1}^{m-1} \sum_{h=1}^{N-1} \binom{N-1}{h} (\alpha_{m'} - \alpha_{m'-1})^h (\alpha_{m'-1})^{N-1-h} \\
 & \quad \times V(s_i, \min\{n-1, h\}, [\alpha_{m'-1}, \alpha_{m'}]) \\
 & + \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-1-j} \min\left\{1, \frac{n}{j+1}\right\} \\
 & \quad \times V(s_i, \min\{n-1, j\}, [\alpha_{m-1}, \alpha_m]) \\
 & + \sum_{k=1}^{n-1} \sum_{j=0}^{N-2} \binom{N-1}{k} \binom{N-1-k}{j} (1 - \alpha_m)^k (\alpha_m - \alpha_{m-1})^j \\
 & \quad \times (\alpha_{m-1})^{N-1-k-j} \min\left\{1, \frac{n-k}{1+j}\right\} V(s_i, k, [\alpha_m, 1]).
 \end{aligned}$$

This expression groups the profiles of opponent messages under which bidder i advances into four terms depending on the highest message sent by her opponents:

- The first term covers profiles in which bidder i is the only buyer to opt-in. She advances for sure, pays c , and gets the asset at a price of 0.
- The second term covers events in which the highest opt-in message sent by any of i 's opponents, m' , is less than m , and the number of opponents sending this message is h . In this case, bidder i advances for sure. Her expected payoff conditional on advancing depends upon the number of opponents who send message m' and advance; this number is h if $h < n - 1$ and $n - 1$ if $h \geq n - 1$.
- The third term covers events in which the highest opt-in message sent by i 's opponents is m , and the number of opponents sending m is j . Bidder i advances for sure if $j < n$, and with probability $n/(j + 1)$ otherwise. Conditional on advancing, the number of opponents determining bidder i 's payoff is either j or $n - 1$, whichever is lower.
- The fourth term covers scenarios in which $k < n$ of bidder i 's opponents send messages higher than m , and j more opponents send message m . If $k + j < n$, then bidder i advances for sure; if $k + j \geq n$, the k bidders who sent messages above m advance for sure, and bidder i advances with probability $(n - k)/(j + 1)$. If she does advance, her payoff is determined by the highest type among the k bidders sending messages above m , whose types are drawn from the interval $[\alpha_m, 1]$.

When $m = 1$, the second term in equation (1) vanishes and, when $m = M$, the last term vanishes. Finally, the payoff to opting out, $v_\alpha(0; s_i)$, is 0.

Equation (1) shows that $v_\alpha(m; s_i)$ depends on α_m and the lower thresholds $\{\alpha_0, \alpha_1, \dots, \alpha_{m-1}\}$, but not on thresholds higher than α_m . In the event that bidder

i advances to the second round after sending message m , every opponent with type above α_m , and who therefore sends a message higher than m , also advances. Consequently, bidder i can treat these opponents as random draws from the interval $[\alpha_m, 1]$, and the payoff from winning against them does not depend on how this interval is partitioned. This is not true for opponents who send messages less than m . In this case, there is some probability that the opponent with the highest type may not advance if $n - 1$ other opponents send the same message as she does. The likelihood of this event depends on how the interval below α_{m-1} is partitioned. However, since bidder i is certain to advance in this event, the payoff associated with it is the same whether she sends message m or $m + 1$. As a result, it drops out of the difference $v_\alpha(m + 1; s_i) - v_\alpha(m; s_i)$, which then depends only on $(\alpha_{m-1}, \alpha_m, \alpha_{m+1})$.

B. Constructing the Equilibrium

As noted in Lemma 2, in a symmetric equilibrium, bidder i must be indifferent between sending message $m + 1$ and m at $s_i = \alpha_m$ for each $m \in \{0, \dots, M - 1\}$. Thus, a symmetric equilibrium must satisfy the M indifference conditions

$$v_\alpha(m + 1; \alpha_m) - v_\alpha(m; \alpha_m) = 0, \quad m \in \{0, \dots, M - 1\}.$$

As noted above, each indifference condition depends only on α_{m-1} , α_m , and α_{m+1} (or on α_0 and α_1 for the indifference condition $v_\alpha(1; \alpha_0) = v_\alpha(0; \alpha_0) = 0$). Further, the difference $v_\alpha(m + 1; \alpha_m) - v_\alpha(m; \alpha_m)$ satisfies a strict single-crossing property in α_{m-1} ; so for a given choice of α_m and α_{m+1} , there is a unique value of α_{m-1} that satisfies $v_\alpha(m + 1; \alpha_m) - v_\alpha(m; \alpha_m) = 0$. We exploit this fact to construct the symmetric equilibrium “from the top down.”

Postponing (for now) the problem of calculating M^* , given Theorem 1, we expect to find an equilibrium with $M = \bar{M}$ if $\bar{M} < M^*$, and with $M = M^*$ otherwise. (In the former case, the second condition in Lemma 2 is automatically satisfied, so a solution to the M indifference conditions constitutes an equilibrium; in the latter case, we will still need to show separately that $v_\alpha(M + 1; 1) \leq v_\alpha(M; 1)$.) Given this value of M , for $t \in [0, 1]$, define $\alpha_M(t) = 1$ and $\alpha_{M-1}(t) = 1 - t$. Define $\alpha_{M-2}(t)$ as the unique value of α_{M-2} satisfying $v_\alpha(M; \alpha_{M-1}(t)) = v_\alpha(M - 1; \alpha_{M-1}(t))$ given $\alpha_M = 1$ and $\alpha_{M-1} = 1 - t$; as noted above, this value is uniquely defined. Similarly, define $\alpha_{M-3}(t)$ to satisfy $v_\alpha(M - 1; \alpha_{M-2}(t)) = v_\alpha(M - 2; \alpha_{M-2}(t))$ given the values of $\alpha_{M-1}(t)$ and $\alpha_{M-2}(t)$, and so on, until $\alpha_0(t)$ is defined by $v_\alpha(2; \alpha_1(t)) = v_\alpha(1; \alpha_1(t))$.

Now, letting $\alpha(t) = (\alpha_0(t), \alpha_1(t), \dots, \alpha_{M-1}(t))$, note that at every value of t , $\alpha(t)$ satisfies the “top” $M - 1$ indifference conditions by construction. Thus, if there is a value t^* of t at which the bottom indifference condition, $v_{\alpha(t^*)}(1; \alpha_0(t^*)) = v_{\alpha(t^*)}(0; \alpha_0(t^*)) = 0$, is satisfied, then the thresholds $\alpha(t^*)$ will be an equilibrium. The proof of existence, then, involves showing that $v_{\alpha(t)}(1; \alpha_0(t))$ is positive at $t = 0$, negative at large t , and continuous in between. Thus, a solution t^* exists to solve $v_{\alpha(t)}(1; \alpha_0(t)) = 0$, and therefore to satisfy the M indifference conditions.

The only remaining issue is the value of M^* . When $M = M^* < \bar{M}$, it’s necessary for equilibrium that bidders with the highest type $S_i = 1$ be unwilling to

deviate from message M (the highest message being used in equilibrium) to higher, unused messages, in order to advance to the auction for certain. By sending message M , bidder i is already assured of advancing except when at least n other bidders send message M as well, in which case every bidder advancing will have a type above α_{M-1} . The condition for unused messages to be unprofitable deviations, then,

$$v_\alpha(M + 1; 1) \leq v_\alpha(M; 1)$$

turns out to be equivalent to the condition

$$V(1, n - 1, [\alpha_{M-1}, 1]) \leq 0$$

i.e., that a bidder with type $s_i = 1$ does not make money, on average, from an auction fully stocked with competitors with types who send message M in equilibrium. Since the left-hand side of this latter expression is decreasing in α_{M-1} , this in turn requires that α_{M-1} be sufficiently close to 1. (If the interval $[\alpha_{M-1}, 1]$ were large, a bidder with type $S_i = 1$ would want to advance even if all her opponents were in this interval; if it is sufficiently small, however, she would not.) As M gets larger, the interval $[\alpha_{M-1}, 1]$ gets narrower; by defining M^* as the highest value of M for which the construction above works, this condition ends up being automatically satisfied for $M = M^*$.¹⁹

It is worth noting that this final necessary condition for equilibrium when $\bar{M} > M^*$ is analogous to the No Incentive To Separate (NITS) condition introduced by Chen et al. (2008).²⁰ In their model, it is an equilibrium refinement, and selects a unique equilibrium, the one with the maximal number of messages. In our setting, since we assume the seller is committed to advancing bidders monotonically based on messages, this is not a refinement, but a condition that must be satisfied in any equilibrium where some messages are not used.

C. Features of Equilibrium

For a specific example of what the equilibrium looks like, we let $N = 5$; let $V_i = S_i$ be drawn uniformly from the interval $[0, 100]$ rather than $[0, 1]$ (so the results are easier to read); and let $c = 5$. Table 1 reports the equilibrium thresholds, revenues, bidder and total surplus for various values of \bar{M} when $n = 2$. In this example, $M^* = 3$, so when $\bar{M} \geq 3$, the equilibrium uses only messages $\{0, 1, 2, 3\}$. (If we tried to construct an equilibrium with more messages, we would fail because all the thresholds could not fit into the type space $[0, 100]$.) When $\bar{M} < 3$, bidders use all of the available messages. In these cases, the highest type would like to separate but cannot do so because of the constraint that the seller imposes on the size of the

¹⁹More formally, we define (in the proof of Theorem 1) M^* as the highest value of M for which $\alpha_0 < \alpha_1 < \dots < \alpha_{M-2} < \alpha_{M-1}$ exist such that $\alpha_{M-1} = \alpha_M = 1$; the indifference condition $v_\alpha(m + 1; \alpha_m) = v_\alpha(m; \alpha_m)$ holds at each $m = 1, 2, \dots, M - 1$; and $v_\alpha(1; \alpha_0) > 0$.

²⁰In their setting, NITS is the condition that the *lowest*-type sender would not choose to reveal his type truthfully if he could, while in our model, the condition is on the highest type. As noted above, the direction in which the sender/bidder would like to misrepresent his type goes in opposite directions in the two models.

TABLE 1—EQUILIBRIUM PARTITION AND PAYOFFS:
 $N = 5, V_i = S_i \sim U[0, 100], c = 5, n = 2$

Opt-in messages available (\bar{M}):	1	2	3+
α_3	—	—	100.00
α_2	—	100.00	98.12
α_1	100.00	83.79	83.64
α_0	51.50	49.45	49.42
Revenue	53.67	57.21	57.26
Bidder surplus	16.96	15.44	15.42
Total surplus	70.63	72.65	72.68

message space. Note that the intervals are narrower at higher messages. In this sense, there is finer sorting at the top of the type space than at the bottom.²¹

As \bar{M} increases, the opt-in threshold α_0 decreases. As a result, both the probability that a bidder opts in and the expected number of bidders opting in increases with \bar{M} . In addition, as \bar{M} increases and more messages are used, the bidders sort more effectively, and the selected bidders are more likely to be those with the highest types. Both of these effects favor higher revenue: greater participation means that the seller is more likely to sell the asset at a positive price (since this requires at least two bidders to opt in); and better selection implies that the seller is likely to sell for a higher price. Note, however, that bidder surplus goes the opposite direction, decreasing as \bar{M} increases. Bidders benefit heavily from being the only one to opt in, which is more likely when \bar{M} is lower; and they benefit from less effective sorting, since it increases the chance they do not face the toughest possible competition. Still, the increase in revenue appears to dominate the decrease in bidder surplus: in every example we've solved, total surplus is increasing in \bar{M} .²² Thus, when choosing how much to restrict the message space, the seller does not face any trade-off between revenue and efficiency: both argue against restricting the set of messages more than necessary.²³

IV. Welfare and Revenue

In this section, we evaluate the performance of the indicative bidding mechanism. The benchmark we compare it to is an auction with unrestricted entry, where bidders choose (independently and simultaneously) whether to enter. The timing is the same as in our model: bidders learn their types S_i , decide simultaneously whether to enter,

²¹Under the two special cases noted above in which uniqueness holds, this property holds as well when the number of available messages is not a binding constraint, i.e., when $\bar{M} \geq M^*$.

²²We found the same result in numerical examples with private value updating, but these were limited to the case where $n = 2$ and $V_i = S_i + T_i$, with T_i independently distributed according to an exponential distribution.

²³We can similarly examine how the equilibrium partition, and expected payoffs, vary with the number of bidders n advancing to the second round. Here, even the interests of the seller and the buyers are aligned: in the examples we have examined, both expected revenue and bidder surplus are decreasing in n . That is, both the seller and the bidders benefit from reducing the duplication of entry costs, and therefore setting $n = 2$ is optimal from both sides' point of view. Without the "small rents" assumption, however, this would not necessarily be the case, as we discuss further in Section V.

and then those who chose to enter incur the cost of due diligence, learn their values V_i , and submit binding bids.²⁴ We continue to assume a second-price auction, so bidding truthfully remains a dominant strategy for those who enter. The symmetric equilibrium involves a cutoff strategy, in which bidders enter when their type S_i is above some threshold γ , which we refer to as the entry threshold; the number of entrants is therefore a binomially distributed random variable.

The comparison between the indicative bidding mechanism and a standard auction with endogenous entry involves a clear trade-off. On the one hand, by limiting the number of bidders who perform due diligence and make binding bids, the indicative bidding mechanism introduces the risk that the seller will not receive bids from the “right” bidders—those with the two highest valuations among those who opted in. This could happen due to tiebreaking eliminating a bidder with one of the two highest initial signals, or due to a bidder with a lower initial signal having an unexpectedly high valuation. On the other hand, because the indicative bidding mechanism reduces the “risk” to a marginal bidder who opts in—in the event competition is very strong, she won’t be selected and won’t incur the entry cost—it reduces the threshold at which bidders are willing to opt in, and therefore increases the likelihood that two or more bidders opt in and the seller earns positive revenue.

To see which of these two effects dominates, we use the following example. We let $V_i = S_i + T_i$, where $\{T_i\}$ are new signals learned upon entry and are independent of $\{S_i\}$. We let $c = 5$, $\{S_i\}$ be independently and identically distributed uniformly on $[0, 100]$, and $\{T_i\}$ be independently and identically distributed according to the exponential distribution with parameter $\lambda = 0.12$.²⁵

Table 2 shows expected payoffs for this example, comparing the indicative bidding mechanism (with $n = 2$ and \bar{M} sufficiently large to not bind) to the standard auction with unrestricted entry. As noted above, for every N , α_0 is less than γ , that is, stochastically more bidders opt in under indicative bidding than enter the unrestricted auction. This means that both the probability of a sale and the probability of positive revenue are higher with indicative bidding. However, since entry is capped at two, the expected revenue conditional on at least two bidders entering is lower with indicative bidding. Still, the effect of greater participation appears to consistently dominate the effect of lesser selection: in this example, as well as in every other example we’ve solved, revenue, bidder surplus, and total surplus are all higher under indicative bidding.²⁶

²⁴This is therefore the “selective entry model” considered by Marmer, Shneyerov, and Xu (2013), but without their assumption that (V_i, S_i) are independent across bidders. The model can be thought of as a hybrid between the model of Levin and Smith (1994), in which bidders have no private information at the time of entry, and the model of Samuelson (1985), in which bidders know their valuations prior to entry and do not update at all post-entry.

²⁵The exponential distribution was chosen because $T_i - T_j$ then follows a Laplace distribution, making $u_2(s_i, s_j)$ straightforward to calculate. In this case, $E \max\{0, T_i - T_j\} = \frac{1}{2\lambda} = 4\frac{1}{6} < c$, so the “small rents” assumption holds.

²⁶The dominance of the indicative bidding mechanism depends on the message space being unconstrained. When N is low and $\bar{M} = 1$ —bidders are restricted to just opting in or opting out, rather than sorting themselves more finely through differential indicative bids—the selection effect (the risk of not advancing the right bidders) is exacerbated, and revenue is lower than under unrestricted entry. In all of the numerical examples we’ve run, however, indicative bidding has consistently Pareto-dominated unrestricted entry when $\bar{M} \geq M^*$, i.e., when the number of messages available is not a binding constraint.

TABLE 2—INDICATIVE BIDDING ($n = 2, M = M^*$) VERSUS UNRESTRICTED ENTRY, VARIOUS N

Potential bidders (N)	3	4	5	7	10	20	50	200
<i>Panel A. Indicative bidding</i>								
α_0 (“opt-in” threshold)	29.59	38.49	45.48	55.66	65.40	79.88	90.82	97.40
Revenue	50.50	59.52	65.63	73.42	79.98	88.77	95.20	99.50
Bidder surplus	23.79	19.50	16.63	12.98	9.95	5.96	3.12	1.24
Total surplus	74.29	79.02	82.25	86.40	89.92	94.73	98.32	100.74
<i>Panel B. Unrestricted entry</i>								
γ (entry threshold)	33.56	44.65	52.68	63.33	72.53	84.84	93.39	98.25
Revenue	50.04	58.24	63.58	70.15	75.65	83.22	89.27	93.28
Bidder surplus	22.91	18.40	15.55	12.15	9.39	5.76	2.87	0.80
Total surplus	72.95	76.63	79.12	82.30	85.03	88.98	92.14	94.08

While we cannot prove that this result holds universally, we can prove that it always holds when N is sufficiently large. Intuitively, as N gets large, the risk under indicative bidding of selecting the “wrong” bidders shrinks, as only bidders with types close to 1 opt in, and the participation effect dominates. In fact, it’s relatively easy to characterize limiting behavior in both the indicative bidding mechanism and the unrestricted auction as N grows, making the comparison straightforward when N is sufficiently large. When N is large, any symmetric equilibrium of the indicative bidding game has $M = 1$ (only messages 0 and 1 are used). Further, as N grows, the opt-in threshold α_0 under indicative bidding and the entry threshold γ for the unrestricted auction both approach 1 at rate $1/N$, so the number of bidders opting in (or entering) approaches a Poisson distribution, with all bidders entering having types arbitrarily close to 1. This allows us to easily characterize expected payoffs in the limit for both mechanisms, leading to the following results.²⁷

THEOREM 2: *Fix an environment other than N . For any indicative bidding mechanism (n, \bar{M}) , if N is sufficiently large, expected revenue and total surplus are both strictly higher than in an unrestricted auction.*

Note that as N grows, the optimal reserve price in the auction with unrestricted entry goes to 0, so we are comparing to the “optimal” standard auction. In the case of large N , we can also characterize which indicative bidding mechanisms perform best.

²⁷Table 2 showed that not only did the indicative bidding mechanism outperform the unrestricted auction in expected revenue and total surplus, but also in bidder surplus. While we can’t prove this in full generality even for large N , we can prove it in the special case where V_i is additively separable into independent components learned before and after entry. Specifically, if $V_i = u(S_i) + T_i$, where u is strictly increasing and continuously differentiable and $\{T_i\}$ are independent of $\{S_i\}$ (but not necessarily independent across i), then for any indicative bidding mechanism (n, \bar{M}) , if N is sufficiently large, bidder surplus is strictly higher than in an unrestricted auction. The analogue to Theorem 3 holds as well: if V_i is additively separable and N is sufficiently large, bidder surplus is strictly decreasing in n , strictly decreases if a reserve price is used, and strictly increases if a small entry subsidy is used.

THEOREM 3: Fix an environment other than N . If N is sufficiently large ...

- (i) The set of messages \overline{M} doesn't matter, as $M = 1$ in equilibrium.
- (ii) Expected revenue and total surplus are both strictly decreasing in n .
- (iii) For any given indicative bidding mechanism (n, \overline{M}) ,
 - A positive reserve price would strictly reduce both expected revenue and total surplus.
 - A small entry subsidy would strictly increase both expected revenue and total surplus.

Proofs of Theorems 2 and 3 are given in the Appendix, but much of the intuition can be gained rather quickly. In both the indicative bidding mechanism and the unrestricted auction, as N grows, total bidder surplus goes to 0,²⁸ so showing that total surplus is higher with indicative bidding, or decreases with n , or decreases with a positive reserve price, also establishes the same properties for revenue.

The key to all of these results is that participation in the indicative bidding mechanism, while higher than participation in the unrestricted auction, is still below the socially optimal level. This is shown formally in the Appendix, but can best be understood in the following way. In a standard English auction with private values, one bidder's entry decision imposes no net externality on the rest of the environment. (If the new entrant wins, her payoff is the difference between her own valuation and the second-highest, which is exactly her contribution to social surplus. If she loses and sets the price, her presence transfers surplus from the winner to the seller; if she loses and doesn't set the price, her presence has no effect.) However, in the indicative bidding mechanism, a bidder's entry decision can impose an externality, because her choice to opt in might lead to another bidder not advancing to the auction. This happens only when at least n bidders are already opting in, and therefore the auction would be "full" regardless. When N is large, all entrants have roughly the same type, so a marginal entrant's entry decision (when at least n others are opting in) does not affect the seller's or the other entrants' expected payoffs; the bidder who would have advanced otherwise but now does not, by the small rents assumption, was anticipating a negative payoff from the auction, and so the net externality caused by the entrant is *positive*. As a result, participation is below the socially optimal level, and any change that increases participation—lowering n , reducing the reserve price, or subsidizing entry—increases total surplus on the margin via participation. And none of these changes lower total surplus directly; so each of them increases surplus overall, and therefore increases revenue as well.

²⁸ Expected bidder surplus must be 0 at type $S_i = \alpha_0$ or $S_i = \gamma$, so if the effect of S_i on the distribution of V_i is bounded, a single bidder's ex ante expected surplus can be thought of as the area of a triangle whose base and height are both proportional to $1/N$, so the N bidders' combined surplus is proportional to $N/N^2 = 1/N$.

V. Extensions

A. Pure Cheap Talk

As is usual in the mechanism design literature, we have assumed the seller can credibly commit to the mechanism used, which includes how he responds to each message. Thus, the messages are not true cheap talk, as their “meaning” is built into the mechanism.

However, we can easily reverse this assumption. Suppose that bidders have a true “opt out” option, but beyond that, the indicative bids they submit (or the messages they send) are not binding on the seller in any way; the seller is committed to advancing a maximum of n bidders to the auction (and is committed to an English auction with no reserve price), but can select whichever bidders he wants out of those who opted in.

In this case (assuming \bar{M} is large), a symmetric equilibrium exists with M opt-in messages used, for each $M = 1, 2, 3, \dots, M^*$. Returning to Lemma 2, the second condition for equilibrium no longer needs to hold, since pessimistic off-equilibrium-path beliefs by the seller can be used to deter deviations away from any smaller set of equilibrium messages. Thus, for any M , we can find the thresholds $(\alpha_0, \dots, \alpha_{M-1})$ satisfying the $M - 1$ indifference conditions, and use those to construct equilibrium strategies. Thus, if indicative bids are pure cheap talk, the equilibrium we found earlier is still an equilibrium, but we are now subject to the usual multiplicity problem common to cheap talk games.

B. Relaxing “Small Rents”

We have focused on the case of “small rents,” imposing the assumption that $u_2(s_i, s_j) < c$ for $s_j \geq s_i$. This is a key step in showing that $v_\tau(m + 1; \alpha_m) - v_\tau(m; \alpha_m)$ is single-crossing in α_{m-1} , allowing us to construct equilibrium thresholds “down from the top” of the type space.

Based on numerical examples, “small” violations of this inequality do not have a discontinuous effect on equilibrium. But without this assumption, we cannot in general prove existence of a symmetric equilibrium. However, if we focus on the case of large N , we can characterize what must happen in any symmetric equilibrium, should one exist.

Let $u_n(1, \mathbf{1})$ denote the expected auction payoff to a bidder with type $S_i = 1$ facing $n - 1$ opponents with type 1 as well. Since $u_n(1, \mathbf{1})$ is decreasing in n , define n^* as the unique solution to

$$u_{n^*}(1, \mathbf{1}) - c > 0 \geq u_{n^*+1}(1, \mathbf{1}) - c.$$

Thus, n^* is the largest auction in which a bidder with type $S_i = 1$ would willingly compete, even if all her competitors were as strong as she. (The small rents assumption would imply that $n^* = 1$.) Consider the generic case where both inequalities hold strictly.

If $n > n^*$, then as N gets large, equilibrium is similar to the “small rents” case already analyzed. Very strong bidders opt in, but in the hope that the auction will not be “full”; conditional on $n - 1$ other bidders entering, they would anticipate a negative expected payoff. Only messages 0 and 1 are used in equilibrium, and the opt-in threshold α_0 goes to 1 at rate $1/N$. Above n^* , for the same reasons as in Theorem 3 above, revenue and total surplus are decreasing in n .

If $n \leq n^*$, however, things change. A bidder with type $S_i \approx 1$ anticipates a positive expected payoff from advancing to the auction, even if the auction is “full” and even if all her competitors are also strong. Thus, if $\bar{M} = \infty$, no symmetric equilibrium can exist: the highest type of bidder would always want to deviate to a higher message. If \bar{M} is finite, however, symmetric equilibrium may exist. Let ξ be the solution to

$$V(\xi, n - 1, [\xi, 1]) = 0$$

so that a bidder with type $S_i = \xi$ is exactly indifferent to entering an auction against $n - 1$ opponents with types above ξ . In the limit as N gets large, in any symmetric equilibrium, all bidders with types in the interval $[\xi, 1]$ pool on the highest available message \bar{M} . While we can't tell exactly what the lower types are doing, it doesn't matter for payoffs, as when N is sufficiently large, more than n bidders will have types above ξ with probability going to 1. Thus, there will be no chance of the object not selling; the only inefficiency comes from the fact that with ξ bounded away from 1, the bidders advancing will not be the strongest ones. For this reason, increasing ξ increases both revenue and total surplus, by improving selection; and as a result of this, below n^* , revenue and total surplus are *increasing* in n .

Thus, assuming symmetric equilibrium always exists, both revenue and total surplus are single-peaked in n , with an optimum of either $n = n^*$ or $n = n^* + 1$.

C. First-Price Auctions

Finally, we consider an indicative bidding mechanism where the second-stage auction is a first price sealed-bid auction.

In a first-price auction, a bidder's bid depends both on her own valuation and on her beliefs about her opponents' valuations. This makes the analysis more complicated in a number of ways. While we do not have results about the existence of symmetric equilibrium in our general model, we can construct the equilibrium for the example in Section III (in which $V_i = S_i \sim U[0, 100]$) when $n = 2$, as an illustration of the fact that it may be possible more generally. In this example, a bidder will rationally update her beliefs about her opponents' types based on the new information that she has advanced to the auction, but she does not receive any other information. Thus, there is no loss in imagining that each bidder solves a static problem, choosing both her message m and her bid b at the beginning, just choosing b optimally for the beliefs that will prevail should she advance. As a result, we can treat the game as a static game and use standard mechanism design tools like the envelope theorem to characterize the equilibrium bid function.

A symmetric equilibrium will have the following properties:

- (i) The type space is partitioned into subintervals by thresholds $0 < \alpha_0 < \alpha_1 < \dots < \alpha_{\overline{M}-1} < \alpha_{\overline{M}} = 100$, with types below α_0 opting out and types in the interval (α_{m-1}, α_m) sending message m .
- (ii) On the equilibrium path, those bidders who advance bid $\beta(S_i)$, where $\beta(\alpha_0) = 0$ and β is strictly increasing and continuous on $[\alpha_0, 100]$.

It's not hard to show these are necessary conditions for a symmetric equilibrium. (If there were discontinuities in β or if $\beta(\alpha_0) > 0$, this would mean holes in the support of $\beta(S_i)$, which are impossible in a pay-as-bid auction with a symmetric equilibrium.) Along with these conditions, two more turn out to be necessary and sufficient for equilibrium: indifference of bidders with threshold types α_m between sending message m (and then bidding $\beta(\alpha_m)$ if selected) and sending message $m + 1$ (and then bidding $\beta(\alpha_m)$); and an envelope theorem condition characterizing interim expected payoffs, which constrains the bid function β . (This is formalized as Lemma 4 in the Appendix.)

In the case of second-price auctions, we constructed the equilibrium from the top down: we guessed a value of α_{M-1} , calculated the other thresholds $\alpha_{M-2}, \alpha_{M-3}, \dots, \alpha_0$ required to rationalize it (and each other), and checked whether these thresholds satisfied the "terminal" condition $v_\tau(1; \alpha_0) = 0$, then adjusted the initial value of α_{M-1} until we found thresholds that did. In the case of first-price auctions, we work in the other direction: we will guess a value of α_0 and build the equilibrium up from there, finally checking whether a terminal condition holds at the *top* of the type space. This is because in addition to calculating message thresholds α_m , we also need to construct the equilibrium bid function β , and this is easier from the bottom of the type space.

The equilibrium construction is shown in the Appendix. We should note that we do not have a theoretical proof that this construction will always work. However, we have tried it numerically for various values of c , N , and \overline{M} , and have never yet failed to find an equilibrium.

To facilitate comparisons with the second-price auction, Table 3 illustrates the equilibrium for the numerical example of Section III in which $c = 5$, $N = 5$, and $n = 2$. For each value of \overline{M} , there is a unique symmetric equilibrium in which all of the messages are used. The intervals are narrower at higher messages, so there is finer sorting at the top of the type space, particularly at higher values of \overline{M} . Even though bidders now use all of the available messages, the impact of the additional messages becomes minimal very quickly. The lower thresholds quickly asymptote, as do expected payoffs; the additional messages only serve to divide up the very top of the type space more finely. For example, when $\overline{M} = 10$, the top four messages are used only by bidders with types above 99.44, and 90 percent of bidders who opt in use message 1, 2, or 3. In fact, up to the precision shown in Table 3, revenue and bidder surplus do not change with \overline{M} once \overline{M} is above 4. Thus, even though there is no "natural" upper bound on the number of messages, this property turns out not to be payoff-relevant.

TABLE 3— $N = 5, V_i = S_i \sim U[0, 100], c = 5, n = 2$, FIRST-PRICE AUCTION, VARIOUS \bar{M}

	Opt-in messages available (\bar{M})						
	1	2	3	4	5	7	10
α_{10}	–	–	–	–	–	–	100.00
α_9	–	–	–	–	–	–	99.94
α_8	–	–	–	–	–	–	99.87
α_7	–	–	–	–	–	100.00	99.73
α_6	–	–	–	–	–	99.47	99.44
α_5	–	–	–	–	100.00	98.83	98.83
α_4	–	–	–	100.00	97.71	97.56	97.56
α_3	–	–	100.00	95.21	94.91	94.89	94.89
α_2	–	100.00	89.92	89.27	89.21	89.20	89.20
α_1	100.00	78.47	76.96	76.82	76.81	76.81	76.81
α_0	51.50	48.54	48.26	48.24	48.23	48.23	48.23
Revenue	53.67	57.40	57.69	57.71	57.71	57.71	57.71
Bidder surplus	16.96	15.45	15.33	15.32	15.32	15.32	15.32
Total surplus	70.63	72.84	73.02	73.03	73.03	73.03	73.03

The above example illustrates an important difference between the equilibrium of first-price and second-price mechanisms. A bidder in a first-price auction always shades her bid by more than c below her value and since she pays her bid, she always earns a positive payoff from winning. Consequently, the small rents assumption no longer implies a minimal width for each element of the equilibrium type partition, and there is therefore no natural upper bound on the number of messages used in equilibrium. Furthermore, the bidder with the highest possible type strictly prefers to advance to the auction even if her opponents have types very close to her own. This means that no matter how narrow the top interval gets, a bidder with the highest type would still choose to deviate if an unused high message was available so that she can be selected with probability 1. Hence, all messages are used and an equilibrium is only possible when $\bar{M} < \infty$.

As in the second-price mechanism, as \bar{M} increases, expected revenue increases due to both greater participation and finer selection; bidder surplus decreases; and total surplus increases. Comparing Table 3 to Table 1, revenue and total surplus are marginally higher under the first-price than under the second-price mechanism, but this has nothing to do with the extra messages used, as it holds for \bar{M} as low as 2. Instead, this appears to be driven by the equilibria shown in Table 3 having a slightly lower opt-in threshold α_0 , and having thresholds which are more evenly distributed across the interval $[0, 100]$, leading to both more participation and better sorting. (Consistent with our earlier discussion, these same differences also lead to bidder surplus being lower in the first-price mechanism, but total surplus being higher.)

In the case of private value updating (where new information about V_i is learned during due diligence), equilibrium would be more complicated, because for a given realized valuation V_i , a bidder’s second-stage bid will still depend on the message she sent at the first stage. (This is because she rationally updates her beliefs about her opponents’ types conditional on the event that she herself advanced to the auction, and this updating depends on the message that she herself sent.) Even in the simple example where $V_i = S_i$, a bidder who deviated in the first stage would also

optimally bid differently in the second stage, and we would need to account for “two-stage deviations” like this in verifying that we have found an equilibrium.

VI. Conclusion

We have developed a theory of indicative bidding. The theory establishes that, when entry is costly and expected rents from the private information obtained from entry are small relative to those costs, a seller can use indicative bids to “thin the field” and then hold an auction among a smaller number of buyers. The indicative bids are informative, and their use often leads to greater efficiency and higher revenue than an auction with unrestricted entry, particularly when the number of buyers is large. The theory explains the widespread use of indicative bidding in takeover auctions, where buyers need to conduct costly due diligence prior to submitting binding bids.

The theory also provides an empirical framework for structural estimation of takeover auctions. The goal would be to use observable variation in the number of potential buyers, the number and bids of buyers submitting indicative bids, the number and bids of buyers submitting final bids, and the deal premium across heterogeneous auctions to estimate entry costs and the joint distribution of signals and values of buyers. Estimates of these model primitives could be used to quantify how much information buyers gain from due diligence (allowing one to directly test the “small rents” assumption) and to study how the deal premia would change if the target firm were sold via a negotiation rather than an auction. Gentry and Stroup (2017) develops an estimation strategy to address these issues but, due to limited data at the time, did not use information on indicative bids.²⁹ Our theory suggests that these bids provide useful information for identifying and estimating model primitives and should be incorporated into the econometric model.

MATHEMATICAL APPENDIX

This Appendix contains proofs of Theorems 2 and 3 (the revenue and efficiency properties of the indicative bidding mechanism when N is large), as well as the construction of the equilibrium for a first-price auction with indicative bidding (Section V). Since much of the intuition for equilibrium construction (and therefore existence) for second-price auctions is given in the example in the text, the full proof of Theorem 1 (and the related lemmas) is contained in a separate, online Appendix. The online Appendix also contains a discussion of the results on uniqueness of equilibrium (discussed in Section II following Theorem 1) and proofs of the bidder surplus results for large N (discussed in footnote 27).

A. Preliminaries

To begin, we prove one fact mentioned in the text, which we will make use of multiple times below.

²⁹They (private communication, September 9, 2016) are in the process of using the SEC filings to construct a more detailed dataset on takeover auctions that includes information on indicative bids.

LEMMA 3: *There exists $\varepsilon > 0$ such that $u_2(s_i, s_j) \geq c$ requires $s_i > s_j + \varepsilon$.*

PROOF:

This follows from the assumptions that $u_2(\cdot, \cdot)$ is continuous in both arguments and $u_2(s, s) < c$ for every s . Suppose it were false. Then, for any $\delta > 0$, one could find $s \in [\delta, 1]$ such that $u_2(s, s - \delta) \geq c$. Let $\delta_\ell = 1/\ell$, and let $\{s_\ell\}_{\ell=1,2,\dots}$ be a sequence such that $u_2(s_\ell, s_\ell - \delta_\ell) \geq c$ for every ℓ . Since $s_\ell \in [0, 1]$ are bounded, $\{s_\ell\}_{\ell=1,2,\dots}$ has a convergent subsequence; let $\{s_{j(k)}\}_{k=1,2,\dots}$ be such a subsequence, and let $s_* = \lim s_{j(k)}$. Then $u_2(s_{j(k)}, s_{j(k)} - \delta_{j(k)}) \geq c$ for every k , with $\{s_{j(k)}\} \rightarrow s_*$ and $\{\delta_{j(k)}\} \rightarrow 0$; by continuity of u_2 , $u_2(s_*, s_*) \geq c$, giving a contradiction. ■

B. Characterization of Equilibrium as N Gets Large

For N Large Enough, $M = 1$ in Every Symmetric Equilibrium.—Let ε be the value defined in Lemma 3. Suppose an equilibrium existed with $M > 1$. As we show in the proof of Theorem 1 in the online Appendix, this would require $u_2(\alpha_1, \alpha_0) > c$, which in turn requires that $\alpha_1 - \alpha_0 > \varepsilon$ and therefore $\alpha_0 < 1 - \varepsilon$. Let $-\underline{u} = \max_{s \in [0,1]} \{u_2(s, s) - c\}$ be the largest (i.e., least negative) payoff a bidder can get from entering against an equally strong opponent; and let $\bar{V} = E(V_i | S_i = 1)$.

Consider a bidder with type $S_i = \alpha_0$. If she opts in, she has a α_0^{N-1} chance of being the only one to enter, in which case her payoff will be $E(V_i | S_i = \alpha_0) - c < E(V_i | S_i = 1) - c = \bar{V} - c$; and a probability $(N - 1)\alpha_0^{N-2}(1 - \alpha_0)$ of advancing against one other opponent with type above α_0 and therefore earning a payoff at most $-\underline{u}$. Thus, even ignoring all the scenarios in which multiple other bidders opt in and she may advance and earn a negative expected payoff, our bidder’s payoff from opting in is at most

$$\begin{aligned} v_\tau(1; \alpha_0) &\leq \alpha_0^{N-1}(\bar{V} - c) - (N - 1)\alpha_0^{N-2}(1 - \alpha_0)\underline{u} \\ &\leq \alpha_0^{N-2}[\alpha_0(\bar{V} - c) - (N - 1)\varepsilon\underline{u}] \\ &\leq \alpha_0^{N-2}\varepsilon\underline{u} \left[\frac{\bar{V} - c}{\varepsilon\underline{u}} - (N - 1) \right], \end{aligned}$$

since $\alpha_0 \leq 1$. Thus, for $N > 1 + \frac{\bar{V} - c}{\varepsilon\underline{u}}$, $v_\tau(1; \alpha_0)$ would have to be strictly negative in a symmetric equilibrium with $M > 1$, and therefore no such equilibrium can exist.

As N grows, $\alpha_0 \rightarrow 1$ but $\alpha_0^N \not\rightarrow 0$ or 1.—As noted above,

$$\begin{aligned} v_\tau(1; \alpha_0) &\leq \alpha_0^{N-1}(\bar{V} - c) - (N - 1)\alpha_0^{N-2}(1 - \alpha_0)\underline{u} \\ &\leq \alpha_0^{N-2}[(\bar{V} - c) - (N - 1)(1 - \alpha_0)\underline{u}]. \end{aligned}$$

If α_0 does not go to 1, or converges at a rate slower than $1/N$, then $(N - 1)(1 - \alpha_0) \rightarrow +\infty$, in which case $v_\tau(1; \alpha_0) < 0$, which is impossible.

However,

$$v_\tau(1; \alpha_0) \geq \alpha_0^{N-1}(E(V_i|S_i = \alpha_0) - c) + (1 - \alpha_0^{N-1})(-c),$$

since $E(V_i|S_i = \alpha_0) - c$ is the bidder's payoff if all of her opponents opt out and $-c$ is a lower bound on her payoff if they don't. If α_0 converges to 1 faster than $1/N$, then $(N - 1)(1 - \alpha_0) \rightarrow 0$, which means $\alpha_0^{N-1} \rightarrow 1$. (If the expected number of opponents opting in goes to 0, then the probability that none of them enter goes to 1.) In that case, $v_\tau(1; \alpha_0) \rightarrow E(V_i|S_i = 1) - c > 0$, again a contradiction. So α_0 must go to 1 at rate $1/N$, making α_0^N converge to an interior limit.

This means as N grows, the number of bidders opting in approaches a Poisson random variable, with all entrants having types arbitrarily close to 1. We can calculate the Poisson parameter by noting that, even away from the limit, α_0 satisfies

$$0 = \sum_{k=0}^{n-1} \Pr(k : N - 1) V(\alpha_0, k, [\alpha_0, 1]) + \sum_{k=n}^{N-1} \Pr(k : N - 1) \frac{n}{k+1} V(\alpha_0, n - 1, [\alpha_0, 1]),$$

where $\Pr(k : N - 1)$ is the probability that exactly k of a bidder's $N - 1$ opponents opt in and (as before) $V(\alpha_0, k, [\alpha_0, 1])$ is the expected surplus (net of costs) a bidder with type α_0 earns in an auction with k opponents with types drawn randomly from $[\alpha_0, 1]$. As N grows and $\alpha_0 \rightarrow 1$,

$$V(\alpha_0, k, [\alpha_0, 1]) \rightarrow V(1, k, \{1\}) = u_{k+1}(1, \mathbf{1}) - c,$$

and $\Pr(k : N - 1)$ approaches the probability defined by the Poisson distribution. Let $u_k \equiv u_k(1, \mathbf{1})$; in the limit, the equilibrium participation level λ_n is the solution to

$$0 = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (u_{k+1} - c) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} (u_n - c).$$

To show that this has a unique solution, define the right-hand side as $V_n(\lambda)$, and note that

$$\begin{aligned} \frac{\partial}{\partial \lambda} V_n(\lambda) &= \sum_{k=1}^{n-1} \frac{k \lambda^{k-1} e^{-\lambda}}{k!} (u_{k+1} - c) - \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (u_{k+1} - c) \\ &\quad + \sum_{k=n}^{\infty} \frac{k \lambda^{k-1} e^{-\lambda}}{k!} \frac{n}{k+1} (u_n - c) - \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} (u_n - c) \\ &= \sum_{k'=0}^{n-2} \frac{\lambda^{k'} e^{-\lambda}}{k'!} (u_{k'+2} - c) + \sum_{k'=n-1}^{\infty} \frac{\lambda^{k'} e^{-\lambda}}{k'!} \frac{n}{k'+2} (u_n - c) - V_n(\lambda). \end{aligned}$$

Now, the two sums are negative, since $u_k < c$ for $k > 1$ under small rents; so this says that whenever V_n is positive, it's strictly decreasing in λ , so it's strictly single-crossing from above; so if a solution exists, it's unique. Thus, $V_n(0) = u_1 - c > 0$, and it's not hard to show that V_n is negative when λ gets large, so a solution does indeed exist.

Limit Welfare.—When $M = 1$, total surplus generated in equilibrium can be written as

$$W(n, N) = \sum_{k=0}^{n-1} \Pr(k : N)(w_k(\alpha_0) - kc) + \sum_{k=n}^N \Pr(k : N)(w_n(\alpha_0) - nc),$$

where $\Pr(k : N)$ is the probability that exactly k out of N bidders opt in and $w_k(\alpha_0)$ is the expected total surplus (gross of entry costs) generated by an auction with k bidders with types $s_i \in [\alpha_0, 1]$. If we let

$$w_k = \lim_{\alpha_0 \rightarrow 1} w_k(\alpha_0) = E\left\{ \max\{V_1, \dots, V_k\} \mid S_1 = \dots = S_k = 1 \right\},$$

then as N gets large and $\alpha_0 \rightarrow 1$, $w_k(\alpha_0) \rightarrow w_k$. As N gets large, $\Pr(k : N)$ goes to the Poisson distribution, so

$$W_n \equiv \lim_{N \rightarrow \infty} W(n, N) = \sum_{k=0}^{n-1} \frac{\lambda_n^k e^{-\lambda_n}}{k!} (w_k - kc) + \sum_{k=n}^{\infty} \frac{\lambda_n^k e^{-\lambda_n}}{k!} (w_n - nc),$$

where λ_n is the equilibrium Poisson parameter.

Now, fixing n , let

$$W_n(\lambda) \equiv \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (w_k - kc) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} (w_n - nc)$$

so that $W_n = W_n(\lambda_n)$ but $W_n(\cdot)$ is also defined at nonequilibrium levels of participation. Next, we calculate

$$\begin{aligned} \frac{\partial}{\partial \lambda} W_n(\lambda) &= \sum_{k=1}^{n-1} \frac{k \lambda^{k-1} e^{-\lambda}}{k!} (w_k - kc) - \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (w_k - kc) \\ &\quad + \sum_{k=n}^{\infty} \frac{k \lambda^{k-1} e^{-\lambda}}{k!} (w_n - nc) - \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} (w_n - nc) \\ &= \sum_{k'=0}^{n-2} \frac{\lambda^{k'} e^{-\lambda}}{k'!} (w_{k'+1} - (k' + 1)c) - \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (w_k - kc) \\ &\quad + \sum_{k'=n-1}^{\infty} \frac{\lambda^{k'} e^{-\lambda}}{k'!} (w_n - nc) - \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} (w_n - nc) \end{aligned}$$

$\left(\frac{k \lambda^{k-1} e^{-\lambda}}{k!} = \frac{\lambda^{k-1} e^{-\lambda}}{(k-1)!}\right)$; the first and third sums then simply substitute k' for $k - 1$, changing the range over which the sum is taken accordingly). Simplifying, then,

$$\frac{\partial}{\partial \lambda} W_n(\lambda) = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (w_{k+1} - w_k - c).$$

The next cool thing to notice: $w_{k+1} - w_k = u_{k+1}!$ This is because in a private-values auction, a single bidder's expected payoff is exactly her contribution to total surplus—the expected difference between the highest valuation when she's there and the highest valuation when she's not. So,

$$\frac{\partial}{\partial \lambda} W_n(\lambda) = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (u_{k+1} - c).$$

Note that this is exactly the first term in the expression for $V_n(\lambda)$, and that the second term in $V_n(\lambda)$ is negative, so $\frac{\partial}{\partial \lambda} W_n(\lambda) > V_n(\lambda)$. But we know that for $\lambda \leq \lambda_n$, $V_n(\lambda) \geq 0$, and therefore $\frac{\partial}{\partial \lambda} W_n(\lambda) > 0$. So for participation levels at and below the equilibrium level, W is strictly increasing in participation.

Limit Bidder Surplus Is 0.—Since λ_n is (by virtue of being the Poisson parameter) the expected number of bidders opting in, if we write α_0 as $\alpha_0(N)$ to emphasize its dependence on N , as N gets large,

$$N(1 - \alpha_0(N)) \rightarrow \lambda_n.$$

For N large, then, $1 - \alpha_0(N) \approx \lambda_n/N$, or

$$\alpha_0(N) \approx 1 - \frac{\lambda_n}{N}.$$

Suppressing the suffix on λ and the dependence of α_0 on N , we know a bidder's expected payoff is

$$v_\tau(1; s_i) = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} V(s_i, k, [\alpha_0, 1]) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} V(s_i, n-1, [\alpha_0, 1]).$$

Since $v_\tau(1; \alpha_0) = 0$ in equilibrium, we can subtract that off and write this as

$$v_\tau(1; s_i) = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (V(s_i, k, [\alpha_0, 1]) - V(\alpha_0, k, [\alpha_0, 1])) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} (V(s_i, n-1, [\alpha_0, 1]) - V(\alpha_0, n-1, [\alpha_0, 1]))$$

and therefore

$$v_\tau(1; s_i) \leq \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (V(1, k, \{\alpha_0\}) - V(\alpha_0, k, \{1\})) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} (V(1, n-1, \{\alpha_0\}) - V(\alpha_0, n-1, \{1\})).$$

Now, a single bidder's ex ante expected surplus is

$$u = \int_0^{\alpha_0} v_\tau(0; s_i) ds_i + \int_{\alpha_0}^1 v_\tau(1; s_i) ds_i = \int_{\alpha_0}^1 v_\tau(1; s_i) ds_i.$$

Multiplying by N , then, total bidder surplus is

$$N \cdot u = N \int_{\alpha_0}^1 v_\tau(1; s_i) ds_i \leq N(1 - \alpha_0) \left[\sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (V(1, k, \{\alpha_0\}) - V(\alpha_0, k, \{1\})) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \frac{n}{k+1} (V(1, n-1, \{\alpha_0\}) - V(\alpha_0, n-1, \{1\})) \right].$$

As $N \rightarrow \infty$, $N(1 - \alpha_0) \rightarrow \lambda_n$; and by continuity, as $\alpha_0 \rightarrow 1$, each of the difference terms $V(1, k, \{\alpha_0\}) - V(\alpha_0, k, \{1\})$ in both summands go to 0, so the entire expression in square brackets goes to 0, so total bidder surplus goes to zero.

C. Proof of Theorems 2 and 3

We prove Theorem 3 first. Part (i) ($M = 1$ in any symmetric equilibrium) has already been shown above.

For part (ii), since limit bidder surplus is 0, limit revenue equals limit welfare, so it suffices to show that W_n is strictly decreasing in n . We've shown that as N grows, total surplus approaches a limit $W_n(\lambda_n)$, where

$$W_n(\lambda) = \sum_{k=0}^{n-1} \frac{\lambda^k e^{-\lambda}}{k!} (w_k - kc) + \sum_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} (w_n - nc)$$

and that $W_n(\lambda)$ is strictly increasing in λ on $[0, \lambda_n]$. We also noted earlier that $u_k = w_k - w_{k-1}$, and by the small rents assumption, $u_k < c$ unless $k = 1$, which means $w_k - kc$ is decreasing in k . This means that fixing λ , $W_n(\lambda)$ is decreasing in n , since decreasing n by one replaces a bunch of $w_n - nc$ terms with $w_{n-1} - (n - 1)c$, which is larger (less negative).

Next, note that λ_n is decreasing in n . This is because we can rewrite the expression defining λ_n as

$$0 = \sum_{k=0}^{\infty} \frac{\lambda_n^k e^{-\lambda_n}}{k!} h(k),$$

where

$$h(k) = \min\left\{1, \frac{n}{k+1}\right\} (u_{\min\{n, k+1\}} - c).$$

For $n > 1$, an increase in n weakly decreases all of the $h(k)$ terms, since they're all negative except for $h(0)$ and adding either an opponent, or a greater chance of being selected when opposed, lowers a bidder's expected payoff; we showed earlier that this expression is decreasing in λ until after it becomes negative, so lowering the $h(k)$ terms requires lowering λ to compensate.

All together, these imply that for $n' < n$,

$$W_n = W_n(\lambda_n) < W_{n'}(\lambda_n) < W_{n'}(\lambda_{n'}) = W_{n'},$$

i.e., limit total surplus is strictly decreasing in n . Limit bidder surplus is 0 at every n , so limit revenue equals limit total surplus, and is therefore strictly decreasing in n .

For part (iii), note that adding a reserve price simply lowers a bidder's expected payoff $V(s_i, k, [\alpha_0, 1])$ in each state of the world where she advances, lowering each limit u_k . By the arguments above, this decreases λ . Adding a reserve weakly decreases total surplus independent of λ , since it may prevent sales; and it decreases total surplus via λ , since surplus is increasing in λ at and below the equilibrium level. So adding a reserve lowers limit total surplus; it doesn't change the fact that limit bidder surplus is 0, so it lowers limit revenue as well. A bidder subsidy

works the opposite way; as long as it is small enough that the small rents assumption continues to hold for the modified auction payoffs, and the resulting change in λ is small enough that it remains in the range where welfare is increasing, it increases limit total surplus. Once again, limit bidder surplus remains 0, so limit revenue increases.

Finally, to prove Theorem 2, note that the unrestricted auction case is simply the special case of the indicative bidding game with $n = \infty$, so that the cap on the number of bidders advancing to the second round never binds. This means Theorem 2 follows as a corollary of Theorem 3 part (ii), since “decreasing in n ” implies “higher at finite n than infinite n ” implies “higher with indicative bidding than without.” ■

D. Constructing Equilibrium for the First-Price Auction

Here, we show the construction of symmetric equilibrium in the indicative bidding game when the auction is a first-price auction, $n = 2$, and $V_i = S_i \sim U[0, 1]$. (The example in the text uses $V_i = S_i \sim U[0, 100]$ simply to make the results easier to read.)

First, we establish necessary and sufficient conditions for symmetric equilibrium. We focus on pure strategy equilibria; as in the case of second-price auctions, allowing mixed strategies only expands the set of equilibrium strategies for a measure zero of bidders, and does not change the set of equilibrium payoffs.

A strategy will consist of a mapping

$$\tau : [0, 1] \rightarrow \{0, 1, 2, \dots, \overline{M}\}$$

from types to messages, and a mapping

$$\beta : [0, 1] \rightarrow \mathbb{R}_+$$

from types to second-stage bids. (As noted in the text, a bidder learns nothing between the first and second stage other than whether or not she advanced to the auction, so there is nothing else to condition her bid on and no loss in assuming a bidder chooses her message and bid simultaneously.) We will let

$$v_{\tau, \beta}(m, b; s_i)$$

denote a bidder’s expected payoff given type $S_i = s_i$ if she sends message m and bids b if selected and her opponents are playing the strategy (τ, β) ; equilibrium consists of a strategy (τ, β) such that

$$v_{\tau, \beta}(\tau(s_i), \beta(s_i); s_i) \geq v_{\tau, \beta}(m, b; s_i)$$

for every $s_i \in [0, 1]$, every $m \in \{0, 1, \dots, \overline{M}\}$, and every $b \in \mathbb{R}_+$.

It is straightforward to show that if (τ, β) is a symmetric equilibrium, τ must be weakly increasing; as before, define α_m as the supremum of the set of types playing

messages m or lower with positive probability. We can then write the probability of advancing to the auction, given τ and one's own message m , as

$$L_\tau(m) = (\alpha_{m-1})^{N-1} + \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-1-j} \left(\frac{2}{j+1}\right) + \sum_{j=0}^{N-2} (N-1) \binom{N-2}{j} (1 - \alpha_m) (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-j-2} \left(\frac{1}{j+1}\right).$$

(The first term is the probability all of the other bidders send messages below m . In the second term, j denotes the number of opponents sending message m , with the rest of the opponents sending messages below m , leading to a $2/(j + 1)$ chance of advancing. In the third term, j again denotes the number of opponents sending message m , with one opponent sending a message above m and the rest sending messages below m , giving a $1/(j + 1)$ chance of advancing.)

It is also straightforward to show that β must be strictly increasing, continuous, and satisfy $\beta(\alpha_0) = 0$. (If $\beta(\alpha_0) > 0$ or β were discontinuous, this would correspond to a gap in the support of $\beta(S_i)$, which is impossible in a symmetric equilibrium of a pay-as-you-bid auction.) Given these properties, for $b \in [\beta(\alpha_{m-1}), \beta(\alpha_m)]$, we can write the probability of winning (i.e., the probability of advancing to the second stage and then outbidding one's opponent) by sending message m and then bidding b as

$$Q_{\tau,\beta}(m, b) = (\alpha_{m-1})^{N-1} + \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-1-j} \left(\frac{2}{j+1}\right) \left(\frac{\beta^{-1}(b) - \alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right).$$

The first term is the probability that all of bidder i 's opponents sent messages below m , in which case bidder i will advance for sure and also be the high bidder for sure; the second term is the probability that bidder i will advance against an opponent who also sent message j (whose type is therefore uniformly distributed on $[\alpha_{m-1}, \alpha_m]$) times the probability of winning the auction against such an opponent. (Note that this expression only holds for "equilibrium combinations" of message and bid; for $m \neq \tau(\beta^{-1}(b))$, $Q_{\tau,\beta}(m, b)$ would take a different form.)

If we let $Q(s_i) = Q_{\tau,\beta}(\tau(s_i), \beta(s_i))$ denote a bidder's equilibrium probability of winning the auction, a standard envelope theorem argument establishes that the expected payoff to a bidder with type $S_i = s_i$ must be

$$v_{\tau,\beta}(\tau(s_i), \beta(s_i); s_i) = \int_0^{s_i} Q(s) ds.$$

Since we can also decompose the equilibrium expected payoff as

$$v_{\tau,\beta}(\tau(s_i), \beta(s_i); s_i) = Q(s_i)(s_i - \beta(s_i)) - L(\tau(s_i))c,$$

we can equate these two and find

$$\int_0^{s_i} Q(s) ds = Q(s_i)(s_i - \beta(s_i)) - L(\tau(s_i))c$$

from which we can calculate

$$\beta(s_i) = s_i - \frac{1}{Q(s_i)} \left(\int_0^{s_i} Q(s) ds + L(\tau(s_i))c \right).$$

This gives us the following necessary conditions for equilibrium, which also turn out to be sufficient.

LEMMA 4: Fix an indicative bidding mechanism (n, \overline{M}) with $n = 2$ and $\overline{M} < \infty$. A strategy (τ, β) is a symmetric equilibrium if and only if the following all hold:

- τ is weakly increasing, and characterized by thresholds $0 < \alpha_0 < \alpha_1 < \dots < \alpha_{\overline{M}-1} < 1$;
- $v_{\tau, \beta}(m, \beta(\alpha_m); \alpha_m) = v_{\tau, \beta}(m + 1, \beta(\alpha_m); \alpha_m)$ for each $m = 0, 1, \dots, \overline{M} - 1$;
- β is strictly increasing and continuous and $\beta(\alpha_0) = 0$;
- $\beta(s_i) = s_i - \frac{1}{Q(s_i)} \left(\int_0^{s_i} Q(s) ds + L(\tau(s_i))c \right)$.

The arguments for necessity have already been discussed. For sufficiency, as is usual in mechanism design problems, the envelope condition defining β is equivalent to incentive compatibility, which rules out profitable deviations to other types' equilibrium strategies. Ruling out deviations to off-equilibrium combinations of (m, b) is tedious but mechanical. A full proof is available upon request.

Finally, we can construct a strategy (τ, β) satisfying the conditions of Lemma 4 as follows. (We suppress the dependence of $v_{\tau, \beta}(m, b; s_i)$, $Q_{\tau, \beta}(m, b)$, and $L_{\tau}(m)$ on (τ, β) .):

- (i) *Guess a value of α_0 .* Since a bidder can do no better than $V_i - c = S_i - c$ by entering, α_0 must be greater than c , so we start with a value of $\alpha_0 \in [c, 1]$.
- (ii) *Calculate the required value of α_1 .* Since $v(0, \beta(\alpha_0); \alpha_0) = v(1, \beta(\alpha_0); \alpha_0)$ and $\beta(\alpha_0) = 0$, we need

$$0 = (\alpha_0 - 0)Q(1, 0) - L(1)c.$$

Now, $Q(1, 0) = \alpha_0^{N-1}$, since a bidder bidding 0 in the second stage will only win the auction if all her opponents opt out; this means

$$0 = \alpha_0^N - L(1)c.$$

While it's not immediately obvious, $L(1)$ depends only on α_0 and α_1 and is strictly increasing in α_1 , and this therefore uniquely pins down the value of α_1 given the value of α_0 .

- (iii) Calculate the bid function on $[\alpha_0, \alpha_1]$. Once α_0 and α_1 are known, we can calculate $Q(s)$ on $[\alpha_0, \alpha_1]$, and therefore calculate $\beta(s_i)$ on $[\alpha_0, \alpha_1]$ via the envelope theorem condition.
- (iv) Calculate the required value of α_2 . Knowing the value of $\beta(\alpha_1)$, the indifference condition $v(1, \beta(\alpha_1); \alpha_1) = v(2, \beta(\alpha_1); \alpha_1)$ becomes

$$Q(1, \beta(\alpha_1))(\alpha_1 - \beta(\alpha_1)) - L(1)c = Q(2, \beta(\alpha_1))(\alpha_1 - \beta(\alpha_1)) - L(2)c.$$

We already noted that $L(1)$ depends only on α_0 and α_1 , as do $Q(1, \beta(\alpha_1))$ and $Q(2, \beta(\alpha_1))$. (While sending message 2 will make bidder i more likely to advance against an opponent who sent message 2, she will never win the auction in those cases, so $Q(2, \beta(\alpha_1))$ does not depend on α_2 .) In contrast, $L(2)$ does depend on α_2 , and is strictly increasing, so this indifference condition uniquely pins down α_2 given what we already know.

- (v) Calculate the bid function on $[\alpha_1, \alpha_2]$. Once α_2 is known, $Q(s)$ can be calculated for $s \in [\alpha_1, \alpha_2]$; from this, we can calculate $\beta(s_i)$ for $s_i \in [\alpha_1, \alpha_2]$.
- (vi) Continue to iterate in this way. Once $\beta(\alpha_2)$ is known, the indifference condition at α_2 determines α_3 ; once α_3 is known, the envelope condition lets us recover $\beta(s_i)$ for $s_i \in [\alpha_2, \alpha_3]$; and so on.
- (vii) Check the terminal condition at the top. Eventually, the indifference condition at $\alpha_{\bar{M}-1}$ will determine a required value of $\alpha_{\bar{M}}$. If this is equal to 1, we've found an equilibrium; if not, we adjust our initial guess at α_0 and try again.

As we noted in the text, we do not have a proof that this construction will always succeed. (In particular, given $\{\alpha_0, \alpha_1, \dots, \alpha_m\}$ and $\beta(s_i)$ up to α_m , we have not found a way to ensure that a value of α_{m+1} exists satisfying the indifference condition $v(m, \beta(\alpha_m); \alpha_m) = v(m+1, \beta(\alpha_m); \alpha_m)$.) That said, the construction has succeeded for every value of c , N , and \bar{M} we've tried.

REFERENCES

- Athey, Susan, Dominic Coey, and Jonathan Levin. 2013. "Set-Asides and Subsidies in Auctions." *American Economic Journal: Microeconomics* 5 (1): 1–27.
- Athey, Susan, Jonathan Levin, and Enrique Seira. 2011. "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions." *Quarterly Journal of Economics* 126 (1): 207–57.
- Bhattacharya, Vivek, James W. Roberts, and Andrew Sweeting. 2014. "Regulating bidder participation in auctions." *RAND Journal of Economics* 45 (4): 675–704.
- Boone, Audra L., and J. Harold Mulherin. 2007. "How Are Firms Sold?" *Journal of Finance* 62 (2): 847–75.

- Boone, Audra L., and J. Harold Mulherin.** 2009. "Is There One Best Way to Sell a Company? Auctions versus Negotiations and Controlled Sales." *Journal of Applied Corporate Finance* 21 (3): 28–37.
- Bulow, Jeremy, and Paul Klemperer.** 2009. "Why Do Sellers (Usually) Prefer Auctions?" *American Economic Review* 99 (4): 1544–74.
- Burguet, Roberto, and József Sákovics.** 1996. "Reserve Prices without Commitment." *Games and Economic Behavior* 15 (2): 149–64.
- Chen, Ying, Navin Kartik, and Joel Sobel.** 2008. "Selecting Cheap-Talk Equilibria." *Econometrica* 76 (1): 117–36.
- Compte, Olivier, and Philippe Jehiel.** 2007. "Auctions and information acquisition: Sealed bid or dynamic formats?" *RAND Journal of Economics* 38 (2): 355–72.
- Crawford, Vincent P., and Joel Sobel.** 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
- Crémer, Jacques, and Richard P. McLean.** 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56 (6): 1247–57.
- Crémer, Jacques, Yossi Spiegel, and Charles Z. Zheng.** 2009. "Auctions with costly information acquisition." *Economic Theory* 38 (1): 41–72.
- Eső, Péter, and Balázs Szentes.** 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Review of Economic Studies* 74 (3): 705–31.
- Farrell, Joseph, and Robert Gibbons.** 1989. "Cheap talk can matter in bargaining." *Journal of Economic Theory* 48 (1): 221–37.
- Fullerton, Richard L., and R. Preston McAfee.** 1999. "Auctioning Entry into Tournaments." *Journal of Political Economy* 107 (3): 573–605.
- Gentry, Matthew, and Tong Li.** 2014. "Identification in Auctions with Selective Entry." *Econometrica* 82 (1): 315–44.
- Gentry, Matthew, and Caleb Stroup.** 2017. "Entry and Competition in Takeover Auctions." ftp://ftp.repec.org/opt/ReDIF/RePEc/dav/wpaper/2017_Entry_and_Competition_in_Takeover_Auctions.pdf.
- Kagel, John, Svetlana Pevnitskaya, and Lixin Ye.** 2008. "Indicative bidding: An experimental analysis." *Games and Economic Behavior* 62 (2): 697–721.
- Levin, Dan, and James Peck.** 2003. "To grab for the market or to bide one's time: A dynamic model of entry." *RAND Journal of Economics* 34 (3): 536–56.
- Levin, Dan, and James L. Smith.** 1994. "Equilibrium in Auctions with Entry." *American Economic Review* 84 (3): 585–99.
- Li, Tong, and Xiaoyong Zheng.** 2009. "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions." *Review of Economic Studies* 76 (4): 1397–1429.
- Li, Tong, and Xiaoyong Zheng.** 2012. "Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions." *Journal of Econometrics* 168 (1): 29–46.
- Lu, Jingfeng, and Lixin Ye.** 2017. "Optimal Two-stage Auctions with Costly Information Acquisition." <http://econ.ohio-state.edu/lixinye/Research/Optimal-two-stage-auction.pdf>.
- Marmar, Vadim, Artyom Shneyerov, and Pai Xu.** 2013. "What model for entry in first-price auctions? A nonparametric approach." *Journal of Econometrics* 176 (1): 46–58.
- Matthews, Steven A., and Andrew Postlewaite.** 1989. "Pre-play communication in two-person sealed-bid double auctions." *Journal of Economic Theory* 48 (1): 238–63.
- McAdams, David.** 2015. "On the benefits of dynamic bidding when participation is costly." *Journal of Economic Theory* 157 (1): 959–72.
- McAfee, R. Preston, and John McMillan.** 1987. "Auctions with entry." *Economics Letters* 23 (4): 343–47.
- Menezes, Flavio, and Paulo Monteiro.** 2000. "Auctions with endogenous participation." *Review of Economic Design* 5 (1): 71–89.
- Roberts, James W.** 2013. "Unobserved heterogeneity and reserve prices in auctions." *RAND Journal of Economics* 44 (4): 712–32.
- Roberts, James W., and Andrew Sweeting.** 2013. "When Should Sellers Use Auctions?" *American Economic Review* 103 (5): 1830–61.
- Samuelson, William F.** 1985. "Competitive bidding with entry costs." *Economics Letters* 17 (1–2): 53–57.
- Subramanian, Guhan.** 2010. *Negotiauctions: New Dealmaking Strategies for a Competitive Marketplace*. New York: W. W. Norton.
- Tan, Guofu.** 1992. "Entry and R&D in procurement contracting." *Journal of Economic Theory* 58 (1): 41–60.
- Ye, Lixin.** 2004. "Optimal auctions with endogenous entry." *Contributions to Theoretical Economics* 4 (1): 1–27.
- Ye, Lixin.** 2007. "Indicative bidding and a theory of two-stage auctions." *Games and Economic Behavior* 58 (1): 181–207.